

How to use LSMEANS statement in SAS (with example)?

Authored by
stats writer

November 19, 2025

RECOMMENDED CITATION

stats writer (2025). *How to use LSMEANS statement in SAS (with example)?*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=96878>

The **LSMEANS** statement in the SAS System is a powerful tool used primarily for estimating and comparing the means of factor levels in complex statistical models. Specifically, the term LSMEANS stands for Least Squares Means, which are essentially the predicted marginal means for each level of a classification variable, adjusted for all other effects in the model.

Unlike simple arithmetic means, LSMEANS are particularly useful when dealing with unbalanced designs--that is, when the sample sizes across groups are unequal--or when interactions are present. By using this statement, researchers can obtain estimates for a linear combination of classification variable levels. A common application of the LSMEANS statement is performing simple comparisons between means, or executing more sophisticated comparisons, such as testing for linear and quadratic trends across ordered levels.

For instance, a command like `LSMEANS age * sex / DIFF;` exemplifies its utility. Here, `age` and `sex` are the classification variables. This command not only provides the estimated marginal means for males and females within each age group but also calculates the differences between these means, facilitating focused statistical inference.

The Necessity of Post Hoc Analysis in ANOVA

In statistical research, the primary method for determining whether there is a significant difference between the means of three or more independent groups is the one-way ANOVA (Analysis of Variance). The ANOVA procedure tests a global null hypothesis, essentially asking: Are all group means equal? If the overall F statistic calculated from the ANOVA table suggests a sufficiently small p-value--typically below a predetermined significance level ($\alpha=0.05$)--we conclude that there is sufficient evidence to reject the global null hypothesis, meaning that at least one group mean differs from the others.

However, rejecting the global null hypothesis does not specify which particular pairs of group means are statistically different. It merely flags the presence of inequality somewhere within the groups. To pinpoint the exact location and magnitude of these differences, researchers must employ a post hoc test. These tests are designed to perform multiple pairwise comparisons while controlling for the increased risk of a Type I error (false positive) that arises from conducting numerous tests simultaneously.

The **LSMEANS** statement in SAS is the primary vehicle for executing a variety of these post hoc tests. By incorporating specific options within the LSMEANS statement--such as `ADJUST=TUKEY`, `ADJUST=BONFERRONI`, or `ADJUST=SCHEFFE`--users can easily implement methods that correct for multiple comparisons, ensuring reliable and accurate conclusions following a significant ANOVA result. The following detailed example illustrates the practical application of the **LSMEANS** statement.

Case Study: Evaluating Study Methods

Consider a hypothetical scenario where a researcher is investigating the effectiveness of three distinct studying methods (Method A, Method B, and Method C) on exam performance. Thirty students are recruited for the study, and they are randomly assigned to use one of the three specified methods to prepare for a standardized exam. The resulting exam scores are collected to determine if the mean performance differs significantly based on the study method employed.

The raw data collected for each student, categorized by their assigned study method, is crucial for the analysis. This dataset represents the foundation upon which the ANOVA and subsequent post hoc analyses will be built. The goal is to statistically verify whether any of the three study methods yield superior or inferior mean scores compared to the others.

The exam results for the students are detailed below. We must first organize this information into a structured format suitable for processing by SAS, which is typically done using the Data Step.

Method A	Method B	Method C
78	81	84
81	83	88
82	83	88
82	85	89
85	86	90
88	88	93
88	90	95
90	91	98

Step 1: Data Creation and Preparation in SAS

To begin the analysis in SAS, we must first define the dataset containing our variables: `Method` (the classification variable, a categorical factor) and `Score` (the dependent variable, a continuous measure). This is achieved using the SAS Data Step, which is essential for reading, preparing, and manipulating data before statistical procedures are applied. The use of `data` lines allows us to input the raw data directly into the SAS session.

It is critical that the `Method` variable is defined as a character variable using the dollar sign (\$) in the `input` statement, as it represents nominal categories rather than numerical values. The following code snippet demonstrates the creation of the dataset named `my_data`, ready for analysis.

```
/*create dataset*/  
data my_data;  
input Method $ Score;  
datalines;  
A 78  
A 81  
A 82  
A 82  
A 85  
A 88  
A 88  
A 90  
B 81  
B 83  
B 83  
B 85  
B 86  
B 88  
B 90  
B 91  
C 84  
C 88  
C 88  
C 89  
C 90  
C 93  
C 95  
C 98  
;  
run;
```

Step 2: Performing the Initial One-Way ANOVA

Once the data is successfully loaded into the SAS environment, the next logical step is to perform the one-way ANOVA using the **PROC ANOVA** procedure. This procedure is specifically designed for balanced designs, but it is often used as a standard starting point for initial variance analysis. The critical elements within the procedure are the **CLASS** statement, which specifies the categorical independent variable (Method), and the **MODEL** statement, which defines the relationship between the dependent variable (Score) and the independent variable.

The output generated by **PROC ANOVA** provides the foundation for our hypothesis testing. Specifically, we look at the F statistic and its associated p-value to determine if the study method has an overall significant effect on the exam score. If the F-test is significant, we then proceed to the detailed pairwise comparisons using LSMEANS.

```
/*perform one-way ANOVA*/
proc ANOVA data=my_data;
class Method;
model Score = Method;
run;
```

Step 3: Interpreting the ANOVA Results

Executing the **PROC ANOVA** code yields a comprehensive analysis of variance table. This table summarizes the sources of variation, degrees of freedom, sums of squares, mean squares, and, crucially, the F ratio and its corresponding probability value. The interpretation of these statistics is paramount in deciding whether to proceed with post hoc testing.

The ANOVA Procedure					
Dependent Variable: Score					
Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	2	175.5833333	87.7916667	5.26	0.0140
Error	21	350.2500000	16.6785714		
Corrected Total	23	525.8333333			

R-Square	Coeff Var	Root MSE	Score Mean
0.333914	4.698685	4.083941	86.91667

Source	DF	Anova SS	Mean Square	F Value	Pr > F
Method	2	175.5833333	87.7916667	5.26	0.0140

Upon reviewing the generated ANOVA table, we can extract the critical statistical measures:

The overall F Value is calculated as **5.26**.

The corresponding p-value (Pr > F) is **0.0140**.

To contextualize these values, we restate the hypotheses being tested by the ANOVA:

H0: All group means (Mean A = Mean B = Mean C) are equal.

HA: At least one group mean is different from the rest.

Since the calculated p-value (**0.0140**) is less than our specified significance level, $\alpha = 0.05$, we confidently reject the null hypothesis. This statistical conclusion confirms that the mean exam score is not equal across all three studying methods. Consequently, we must now determine which specific pairs of methods exhibit statistically significant differences, requiring the use of the **LSMEANS** statement for post hoc comparisons.

Step 4: Implementing LSMEANS for Post Hoc Comparisons

To identify the specific differences among the means, we transition from **PROC ANOVA** to a procedure that handles complex mean comparisons, such as PROC GLIMMIX. While other procedures like PROC GLM can also use LSMEANS, PROC GLIMMIX is robust and widely utilized for generalized linear mixed models, making it an excellent choice for incorporating the LSMEANS statement.

Within PROC GLIMMIX, the **LSMEANS** statement is included, specifying the classification variable (Method) and, crucially, the adjustment method. For this example, we employ Tukey's Honestly Significant Difference (HSD) test by adding the option `ADJUST=TUKEY`. Tukey's method is one of the most common and robust procedures for controlling the family-wise error rate when conducting all possible pairwise comparisons.

```
/*perform Tukey post-hoc comparisons*/  
proc glimmix data=my_data;  
class Method;  
model Score = Method;  
lsmeans Method / adjust=tukey alpha=.05;  
run;
```

Step 5: Interpreting the Tukey HSD Output

The execution of the above SAS code produces extensive output, but the most relevant table for our post hoc analysis is the one summarizing the pairwise mean differences and their adjusted significance levels. This table presents every possible comparison among the three study methods (A vs. B, A vs. C, and B vs. C).

Differences of Method Least Squares Means Adjustment for Multiple Comparisons: Tukey												
Method	_Method	Estimate	Standard Error	DF	t Value	Pr > t	Adj P	Alpha	Lower	Upper	Adj Lower	Adj Upper
A	B	-1.6250	2.0420	21	-0.80	0.4351	0.7096	0.05	-5.8715	2.6215	-6.7719	3.5219
A	C	-6.3750	2.0420	21	-3.12	0.0052	0.0137	0.05	-10.6215	-2.1285	-11.5219	-1.2281
B	C	-4.7500	2.0420	21	-2.33	0.0301	0.0739	0.05	-8.9965	-0.5035	-9.8969	0.3969

We focus our attention on the **Adj P** column, which contains the adjusted p-values that account for the multiple comparisons performed under Tukey's HSD method. A significant difference is indicated if the adjusted p-value is less than the specified $\alpha = 0.05$.

Upon inspection, we find that only one row exhibits an adjusted p-value below the threshold of 0.05: the row comparing the mean difference between Method A and Method C. This is the only statistically significant difference among the three pairs. Specifically, the findings for the A vs. C comparison are:

The difference in LS Means (Method A - Method C) is calculated as **-6.375**. This signifies that students using Method C scored, on average, 6.375 points higher than students using Method A.

The adjusted p-value for this difference is **0.0137**, which is statistically significant.

The adjusted 95% confidence interval for the true difference in mean exam scores between these two groups is . Since this interval does not contain zero, it reinforces the conclusion of a significant difference.

All other pairwise comparisons (A vs. B and B vs. C) yield adjusted p-values greater than 0.05, leading us to conclude that there are no statistically significant differences in mean exam scores between Method A and Method B, nor between Method B and Method C.

Advanced Adjustment Options in LSMEANS

While Tukey's HSD is a popular choice for all pairwise comparisons, the **LSMEANS** statement offers considerable flexibility by supporting a range of other adjustment methods. The selection of the appropriate method depends heavily on the specific research question and the type of comparisons required (e.g., comparing all groups to a control group versus comparing all possible pairs).

Researchers can specify various adjustment methods using the **ADJUST=** option within the **LSMEANS** statement. These adjustments control the Type I error rate in different ways, ranging from conservative to more powerful tests. Understanding these options allows for precise control over the statistical inference process.

Alternative adjustment methods available for post hoc comparisons using **LSMEANS** include:

BON: The Bonferroni adjustment, known for being highly conservative, used when comparing a small number of means.

DUNNETT: Used specifically for comparing all treatment groups against a single control group.

NELSON: An adjustment method often employed in specific clinical trial contexts.

SCHEFFE: A highly versatile and conservative method suitable for testing any linear contrast, not just pairwise comparisons.

SIDAK: A less conservative alternative to Bonferroni, providing slightly more power.

SMM: The Single-Step Max-t method, which is often equivalent to Tukey's HSD for pairwise comparisons.

The choice of adjustment method should be determined a priori based on the study design and the hypotheses formulated, rather than after observing the data.

The following tutorials provide additional information about ANOVA models and post hoc methodologies:

A Guide to Using Post Hoc Tests with ANOVA